

Provided for non-commercial research and education use.
Not for reproduction, distribution or commercial use.



This article was published in an Elsevier journal. The attached copy is furnished to the author for non-commercial research and education use, including for instruction at the author's institution, sharing with colleagues and providing to institution administration.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Discussion

Response to discussants

Roger Cooke

Department of Mathematics, Delft University of Technology, P.O. Box 5031, 2600 GA Delft, Netherlands

Available online 9 February 2008

1. Reply to Simon French

This is a good moment to acknowledge many fruitful discussions with Simon French during the formative period of the classical model. The theory of asymptotically proper scoring rules for average probabilities had been worked out, but we did not really know what to do with the significance level. This theory told us that there should be a positive cut-off level, but of course did not say what that level should be. Simon French suggested choosing this level by optimizing the combined score of the DM. Some practitioners, most notably W. Aspinall, prefer to set this level low enough so that all experts receive positive weight. This is not really consistent with the scoring rule philosophy, but has worked well in practice. To ensure comparability with other results, optimization has been applied to Aspinall's data in the results reported here.

With regard to the social network weights, French's suggestions are worthwhile and indicate that implementation is not as straightforward as might appear at first sight. I would caution that using *all* citations, not just citations from the group of experts, would involve much more work.

2. Reply to Bob Clemen

I am very grateful for Bob Clemen's many thoughtful remarks and for the considerable effort that has gone into producing his results. This is an excellent example of how these data can be used to advance the discussion of various weighting schemes. I break this response down into the following four topics.

2.1. Scoring rules

Theoretically, I have nothing to add to the discussion in [1]. The goals of rewarding individual elicitation and

E-mail address: cooke@rff.org

obtaining statistically accurate and informative combinations are quite different. To illustrate just how different, Table 1 shows the quadratic score (positive sensed on $[-1, 1]$) for two hypothetical experts giving 1000 next day probabilities of rain. The forecasts are thrown into probability bins ranging from 5% to 95%. The scores for each prediction are summed and divided by the total number of predictions.

Both experts are equally informative in the sense that they both attribute 5% probability to 100 next days, etc. Expert 1 is statistically perfectly accurate, whereas expert 2 is wildly inaccurate statistically. Nevertheless, expert 2 receives a better score.

2.2. Inter-quartile ranges

There are many ways to measure informativeness. The Shannon relative information enjoys the advantages mentioned in the paper, which the inter-quartile range or 90% confidence interval does not. But of course, for a specific purpose some other measure may be more suitable.

2.3. Median

The goal is to find statistically accurate and informative distributions, not point estimates. Good Bayesians know that point estimates require loss functions. Nonetheless, some years ago we experimented on scale invariant scoring rules based on the proximity of the median to the realization, in the spirit of Clemen's suggestion. A DM based on such a scoring rule did have statistically better medians, but did not do so well with regard to statistical accuracy or informativeness. There was no significant difference between the equal weight and performance-based DMs with regard to median-realization proximity [2]. The propriety of median-based scoring rules is of course an issue.

Table 1
Two experts assessing next day probability of rain on 1000 days

	Probability of rain next day										Total
	5%	15%	25%	35%	45%	55%	65%	75%	85%	95%	
Expert 1											
Assessed	100	100	100	100	100	100	100	100	100	100	1000
Realized	5	15	25	35	45	55	65	75	85	95	500
Expert 2											
Assessed	100	100	100	100	100	100	100	100	100	100	1000
Realized	0	0	0	0	0	100	100	100	100	100	500

Quadratic score expert 1 = 0.665; quadratic score expert 2 = 0.835.

2.4. Out of sample

This is the most important issue raised by Clemen: does the performance of the performance-weighted DM (PWDM) persist beyond the set of seed variables? Clemen believes that there is no significant difference between the performance weight DM (PWDM) and the equal weight DM (EWDM) outside the variables on which PWDM has been constructed.

As noted above, PWDM does use optimization to remove a degree of freedom in the definition of the classical model. On every study we routinely perform robustness analysis by removing seed variables (and experts) one at a time and re-computing PWDM. It is not uncommon to see the calibration scores of PWDM fluctuate by a factor of 2 or 3 on 10 seed variables (compare Table 4 of Cooke and Goossens, this volume).

Out-of-sample validation involves basing PWDM on an initial set of seed variables, then using this PWDM on other variables and comparing the performance with EWDM on these other variables. This corresponds to the way PWDM is actually used. We can do this by splitting the set of seed variables into two halves, initializing the model on one half and comparing the performance on the other half. Of course, this requires a relatively large number of seed variables. There are 14 studies with at least 16 seed variables. One of these, “TNO dispersion”, eluded conversion to the format of the windows software and cannot currently be read. That leaves 13 studies. Dividing the seed variables into half gives two validation runs, using the first half to predict the second and conversely. Note that the variables on which the PWDM is initialized in these two runs are disjoint. The item weight PWDM could not be computed without writing a new code, so the choice of item versus global weights is denied the PWDM on this exercise.

The data from the 13 studies are shown in Table 2. In 20 of the 26 studies the out-of-sample PWDM out-performs EWDM. The probability of seeing 20 or more “successes” on 26 trials if PWDM were no better than EWDM is 0.0012.

Table 2
26 out-of-sample validation runs; best performer is in bold

Study ^a	DM	Calibration	Information	Combination
TUD disper	e1	0.42	0.646	0.2713
	PW(2)1	0.21	0.8744	0.1836
	e2	0.39	0.7844	0.3059
TUD depos	PW(1)2	0.005	1.525	0.007624
	e1	0.52	1.119	0.5819
	PW(2)1	0.52	1.42	0.7382
Operrisk	e2	0.73	1.324	0.9669
	PW(1)2	0.59	1.374	0.8108
	e1	0.429	0.2793	0.1198
	PW(2)1	0.5337	0.5749	0.3068
Dikering	e2	0.5337	0.3646	0.1946
	PW(1)2	0.185	1.109	0.2053
	e1	0.025	0.7386	0.01846
	PW(2)1	0.4	0.3859	0.1544
Thermbld	e2	0.025	0.7814	0.01954
	PW(1)2	0.05	0.6451	0.03225
	e1	0.07	0.1424	0.009967
	PW(2)1	0.48	0.5527	0.2653
Realest	e2	0.005	0.1424	0.0007119
	PW(1)2	0.07	0.7305	0.05113
	e1	0.05	0.179	0.008948
	PW(2)1	0.33	0.8572	0.2829
EuDis	e2	0.18	0.1676	0.030168
	PW(1)2	0.35	0.6724	0.2353
	e1	0.52	0.9662	0.5024
	PW(2)1	0.52	1.232	0.6408
PIntDos 6 exp. 39 items	e2	0.02	0.749	0.01498
	PW(1)2	0.08	1.204	0.09635
	e1	0.001	1.108	0.0011089
	PW(2)1	0.11	1.038	0.1141
Soil	e2	0.23	0.3262	0.07502
	PW(1)2	0.44	0.6748	0.2969
	e1	0.001	0.3638	0.0003638
	PW(2)1	0.001	0.4135	0.0004135
Gas Environ	e2	0.0001	1.539	0.0001539
	PW(1)2	0.0001	1.551	0.0001559
	e1	0.0001	1.235	0.0001235
	PW(2)1	0.06	2.01	0.1206
	e2	0.72	1.274	0.9171
	PW(1)2	0.73	2.342	1.71

Table 2 (continued)

Study ^a	DM	Calibration	Information	Combination
AOT 6 exp 20 items	e1	0.1	0.2046	0.02046
	PW(2)1	0.1	0.6685	0.06685
	e2	0.5	0.1793	0.08964
	PW(1)2	0.7	0.5799	0.4059
EU WD	e1	0.11	0.6611	0.07272
	PW(2)1	0.0001	2.048	0.0002048
	e2	0.04	0.7983	0.03193
	PW(1)2	0.04	0.7743	0.03097
Estec-2	e1	0.75	0.2427	0.182
	PW(2)1	0.43	0.3623	0.1558
	e2	0.68	0.07269	0.04943
	PW(1)2	0.35	0.1893	0.06627

e1 denotes the EWDM on the first half of the seed variables, e2 denotes EWDM on the second half. PW(2)1 denotes the PWDM constructed on the second half, predicting the first half, and PW(1)2 denotes the PWDM constructed on the first half predicting the second half.

^aPintDos involved 55 seed items, and eight experts, but two experts assessed only a small number of seed variables. The other experts' seed assessments did not wholly overlap; six experts assessed 39 common seed variables used for this exercise. Similarly, AOT was restricted to six experts who assessed 20 common items. The Gas study was split into a corrosion and an environment panel. Many environment experts were also corrosion experts and their corrosion seed assessments were used in the original study. For this exercise only the environment seeds were used for the environment panel. In the Dikering study the multiple measurements from each measuring station were split.

Clemen reports results on 14 validation studies that are somewhat more pessimistic (9 “success” on 14 trials; I checked half of them and verified his numbers). His method involves removing seed variables singly, computing PWDM on the remaining seeds, and using this PWDM to

predict the eliminated seed. On a study with 10 seed variables there are thus 10 *different* PWDMs. Each pair of the 10 DMs shares eight common seeds. The criteria for selecting the 14 studies are not specified. It is difficult to see how all these factors would affect the results. Perhaps the following reasoning partially explains Clemen's less optimistic result: With a small number of seeds, removing one seed favors experts who assessed *that* seed badly and hurts experts who assessed *that* seed well, thus tilting PWDM toward a bad assessment of *that* seed. This happens on *every* seed, thus cumulating the adverse effect on PWDM. This does not happen when *one* PWDM predicts the entire out-of-sample set of seeds. In any case, Clemen's method is not the same as picking *one* PWDM and comparing it on new observations with the EWDM.

It is not obvious that weights derived from a given scoring rule should perform best with respect to that scoring rule. Consider Fig. 5 of Cooke, ElSaadany and Huang, this volume. There we see that the DM formed from likelihood weights does not have better likelihood scores than PWDM and EWDM.

Clemen closes with a fetching metaphor of the index funds (EWDM), which the stock market guru's (PWDM) never seem able to beat. I fetch a different metaphor. An expert viniculturalist (PWDM) can mix different grapes to produce excellent wine, but would you expect to enhance the performance by mixing bottles of wine (EWDM)?

References

- [1] Cooke. Experts in uncertainty. Oxford: Oxford University Press; 1991.
- [2] Rosland G. Expert judgment and performance-testing of decision makers. TU Delft: Department of Statistics, Stochastics and Operational Research; 1996 May.